Bringing Task- and Data-Parallelism to Analysis of Climate Model Output

Robert Jacob, Jayesh Krishna, Xiabing Xu, Sheri Mickelson, Tim Tautges, Mike Wilde, Rob Latham, Ian Foster, Rob Ross, Mark Hereld, Jay Larson (Argonne National Laboratory); Pavel Bochev, Kara Peterson, Mark Taylor (Sandia National Lab.)
Karen Schuchardt, Jian Yin (Pacific Northwest National Lab.); Don Middleton, Mary Haley, David Brown, Richard Brownrigg, Wei Huang, Dennis Shea, Mariana Vertenstein (NCAR), Kwan-Liu Ma, Jinrong Xie (UC-Davis)

Problem:

Climate model output continues to grow in size from both increased resolution and additional fields

- CAM HOMME at 0.125 degrees
  - Single 3D variable: 616 MB
  - Single 2D variable: 25 MB
  - Single history file: 24 GB
  - 1 year of monthly output: 288 GB
  - 100 years of monthly: 28.8 TB
- CSU GCRM 4km horizontal, 100 levels
  - Single 3D variable (cell center): 16 GB
  - Single 3D variable (cell edge): 50.3 GB
  - Single history file 571 GB
  - 1 year of monthly output: 6 TB
  - 100 years of monthly: .6 PB

Climate models are going beyond their traditional structured, rectangular grids. Atmosphere cubed sphere grid (left), Spherical Centroidal Voronoi Tessellation (right)

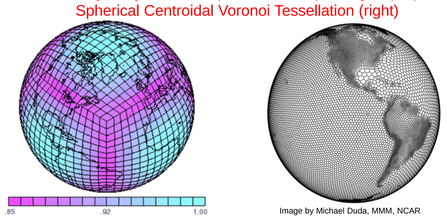

Image by Michael Duda, MMM, NCAR

Analysis packages are often not aware of where data lies on grid.

Location of data on the "CD-grid" of the Community Atmosphere Model Finite Volume dynamical core.

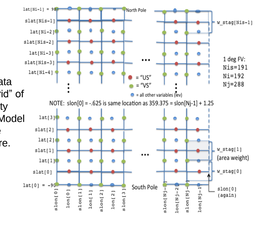

Climate has a very low aspect ratio (10,000km wide while only 10km deep) This means 2D plots are the best way to view climate output. HOWEVER, the current tools to analyze and plot the data are single-threaded and assume rectangular grids!

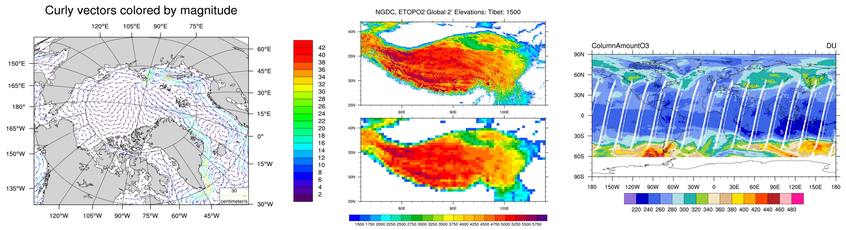

Elements of our solution

NCL (NCAR Command Language) is a widely used scripting language tailored for the analysis and visualization of geoscientfic data.

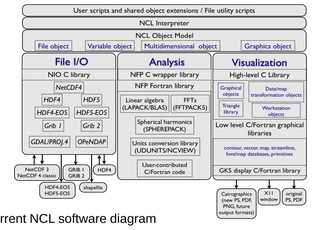

Current NCL software diagram

Mesh-Oriented datABase (MOAB):

- MOAB is a library for representing structur unstructured, and polyhedral meshes, and field data on those meshes.
- Uses array-based storage, for memory efficiency
- Supports MPI-based parallel model

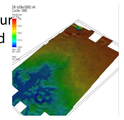

Glacier ice bed with VisIt/MOAB

INteroperable Tools for Rapid dEveloPment of compatIble Discretizations

Intrepid:

- An extensible library for computing operators on discretized fields
- Cell topology, geometry and integration
- Discrete spaces, operators and functionals on cell worksets
- Up to order 10 *H(grad)*, *H(curl)* and *H(div)* FE bases on Quad, Triangle, Tetrahedron, Hexahedron, and Wedge cell topologies
- High quality cubature, e.g., positive weights only on Tri and Tet cells

PNetCDF: NetCDF output with MPI-IO

Based on NetCDF
- Derived from their source code
- API slightly modified
- Final output is indistinguishable from serial NetCDF file

Additional Features
- Noncontiguous I/O in memory using MPI datatypes
- Noncontiguous I/O in file using sub-arrays
- Collective I/O

- **Swift is a parallel scripting system for Grids and clusters**
- **Swift is easy to write**: simple high-level C-like functional language
- **Swift is easy to run**: a Java application. Just need a Java interpreter installed.
- **Swift is fast**: Karajan provides Swift a powerful, efficient, scalable and flexible execution engine.
- Swift is supported by an NSF SI2 grant.

ParVis has increased both the speed and functionality of widely used tools for climate model analysis

Data-parallelism

Using MOAB, Intrepid and PnetCDF, we have reated the Parallel Gridded Analysis Library. ParGALprovides data-parallel and multi-grid capable versions of many typical analysis functions on gridded data sets. ParGAL functions can also operate on the native grids of the output and discretizations of the original model.

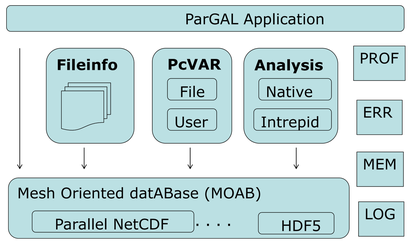

ParGAL Function table (so far)

NCL function Group	NCL Functions	ParGAL Function
File IO	addfile, addfiles	fileinfo, pcvar
Spherical Harmonic Routines	dv2uv* (4 funcs)	divergence
Meteorology	uv2dv_cfd	divergence
Spherical Harmonic Routines	uv2dv* (4 funcs)	divergence
Meteorology	uv2vr_cfd	vorticity
Spherical Harmonic Routines	uv2vr* (4 funcs)	vorticity
Spherical Harmonic Routines	uv2vrdv* (4 funcs)	vorticity, divergence
General Applied Math	dim_avg, dim_avg_n	dim_avg_n
General Applied Math	dim_max, dim_max_n	dim_max_n
General Applied Math	dim_min, dim_min_n	dim_min_n
General Applied Math	max	max
General Applied Math	min	min
Variable Manipulators	delete	pcvar
		gather

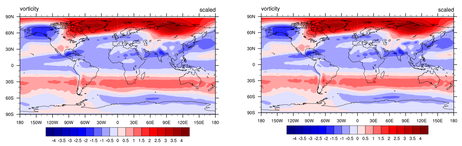

- ParGAL (w/ Intrepid) version of diagnosed vorticity
- Calculated locally on each element
- Easily parallelizable
- Global data not required

- Original NCL version
- Uses spherical harmonics
- Requires global data

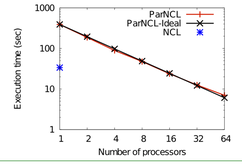

Total execution time for calculating vorticity on a 1 degree grid vs. number of processors for original and ParGAL function called by ParNCL.

ParNCL uses ParGAL, MOAB and NCL graphics to create a parallel version of NCL.

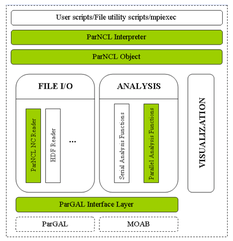

An NCL script executed by ParNCL

```
f = addfiles(diri+fili, "r") ;   open file
tt = f[:]->T(0,{500},{-30:30},:) ;  read a section of data.

wks  = gsn_open_wks("ps","parvis_t")  ; open a PS file
plot = gsn_csm_contour_map(wks, tt(:,:),False)
```

Read with parallel I/O

a parallel data object

In the ParNCL interpreter, tt is gathered to one node and passed to normal NCL graphics routines for plotting.

Task-parallelism

The various working groups of the DOE/NSF Community Earth System Model have each defined a suite of diagnostics that can be applied to the output of each climate model simulation. Hundred of plots are produced by each package. The diagnostic -shell scripts were task-parallelized using Swift.

Ocean Model Diagnostic Package Results:

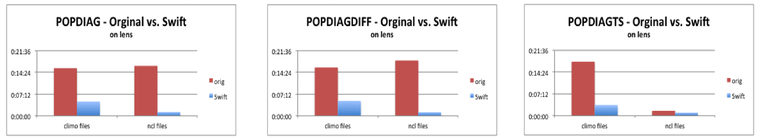

In the above timings, the Ocean Model Diagnostic Package was used to compute the averages of 10 years of 1 degree POP data (climo files) using the NetCDF operators (NCO). NCL was used to create the plots. This was done on 4 compute nodes on lens running a maximum of 8 tasks per node. The original was ran on 1 lens compute node.

Atmospheric Model Diagnostic Package Results:

(a) Time to run the complete diagnostic package on data sets with two different resolutions of the atmosphere model, 0.5 degrees (left) and 0.1 degrees (right; up sampled from 1/8th degree)

(b) Each data set is 12 months long and the diagnostics were performed against observations

(c) Both cases were run on the "midway" cluster of the Research Computing Center at the University of Chicago.

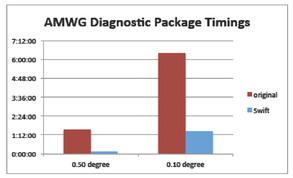

Vision for near-future climate analysis

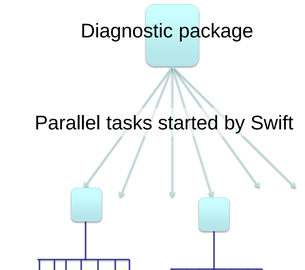

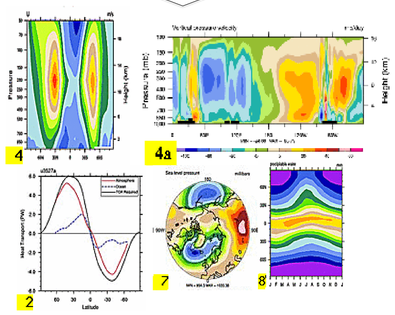

Both the community-developed diagnostic packages and scripts developed by individual scientists can use both task and data-parallelism

Performing averages with a data parallel version of NCO called Pagoda on hopper using 4km data from a GCRM

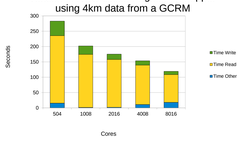

U.S. DEPARTMENT OF ENERGY | Office of Science

Argonne NATIONAL LABORATORY | NCAR | Sandia National Laboratories | Pacific Northwest NATIONAL LABORATORY | UCDAVIS UNIVERSITY OF CALIFORNIA